

SCOR-UQAM NEWSLETTER #1 (October 2023-March 2024)



Arthur Charpentier (UQAM; charpentier.arthur@uqam.ca)

In the era of big data and hardly interpretable machine learning techniques, ensuring fairness and non-discrimination in predictions poses a significant challenge. At the core of the insurance industry lies the differentiation between high-risk and low-risk policyholders. However, with insurers gathering increasingly extensive data, distinguishing between justified pricing based on risk and unfair discrimination based on sensitive factors becomes more intricate. As Kranzberg's law suggests, "*technology is neither inherently good nor bad; nor is it neutral.*" Hence, it's imperative to accurately measure discrimination to prevent its adverse effects. This project aims to address biases within automatic artificial intelligence algorithms utilized for determining optimal pricing in individual policies. Its goal is to mitigate or eliminate biases that could lead to inequities or discriminatory practices based on factors such as gender, race, religion, or origin in the coverage provided by insurers or reinsurers to policyholders. This subject holds both theoretical significance, as it involves enhancing control over the "black box" nature of AI-based models, and practical importance, by reducing the risks associated with discrimination and inequality. In light of the increasing influence of social networks and the consequential reputational challenges faced by insurers and reinsurers, this issue is particularly relevant and timely.

People involved

Several people are involved in the project.

Students (Interns)


Suzie Grondin (alumni: MSc ENSAE), came for six months to work on dynamic pricing games, with reinforcement learning. She also worked on [equipy](#)  




See <https://freakonometrics.hypotheses.org/>

Students (PhD)

Two PhD students started in the Fall 2023, on fairness and discrimination, in insurance markets.


Agathe Fernandes (alumni: MSc IMT Atlantique & Euria), started a PhD at UQAM, supervised by **Arthur Charpentier** and **Ewen Gallic** (AMSE, France), on fairness and calibration of predictive models, 



Olivier Côté (alumni: MSc Laval) started a PhD at Université Laval (Québec), supervised by **Arthur Charpentier** and **Marie-Pier Côté** (Université Laval), on fairness, discrimination and causal models, 




Postdoctoral Fellows

François Hu (alumni: PhD ENSAE-Institut Polytechnique), finished a postdoctoral fellowship, supervised by **Arthur Charpentier** and **Manuel Morales** (Université de Montréal). He will go back to Paris, to join the R&D team of Milliman, 



Two postdoctoral fellowships will be offered, starting during the Summer 2023.

Other People

Ewen Gallic, maître de conférences at Aix-Marseille School of Economics, is invited for one year at UQAM, to work with Arthur Charpentier 



 **Miscellaneous**



December 2023, Christmas Chinese hot pot.

(from the left to the right: Agathe (PhD), Suzie (MSc ENSAE intern), Philipp (PhD), François (postdoc) and Dante (postdoc))




March 2024, “cabane à sucre” (Québec)

(from the left to the right: Agathe (PhD), Dante (postdoc), Noé and Jean (MSc Centrale-Supélec interns), Ewen (invited professor), Hélène (professor))

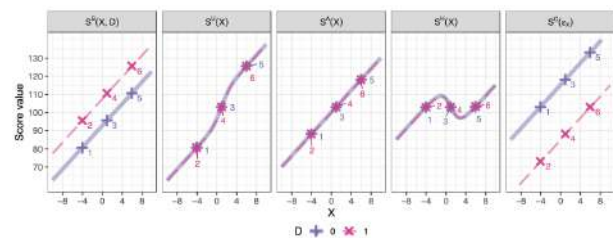


 **Recent Work**

A Fair price to pay: exploiting causal graphs for fairness in insurance

by Olivier Côté, Marie-Pier Côté and Arthur Charpentier,  [SSRN:4709243](https://ssrn.com/abstract=4709243)

In many jurisdictions, insurance companies must not discriminate on some given policyholder characteristics. Omission of prohibited variables from models prevents direct discrimination, but fails to address proxy discrimination, a phenomenon especially prevalent when powerful predictive algorithms are fed with an abundance of acceptable covariates. The lack of formal definition for key fairness concepts, in particular indirect discrimination, hinders the fairness assessment of methodologies. We review causal inference notions and introduce a causal graph tailored for fairness in insurance. Exploiting these, we discuss potential sources of bias, formally define direct and indirect discrimination, and study the properties of fairness methodologies. A novel categorization of fair methodologies into five families (best-estimate, unaware, aware, hyperaware, and corrective) is constructed based on their expected fairness properties. A comprehensive pedagogical example illustrates the practical implications of our findings: the interplay between our fair score families, group fairness criteria, and sources of discrimination.

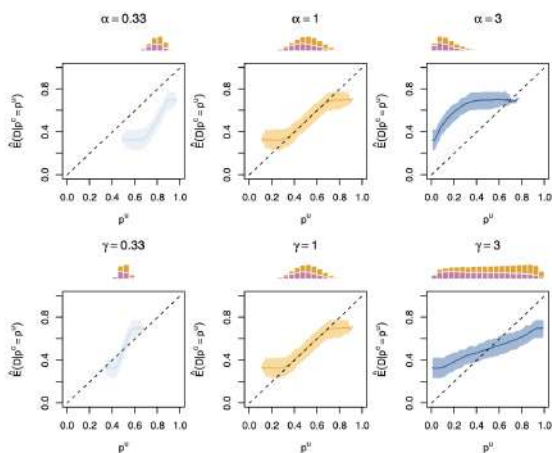


This work has been presented at the **Workshop on Insurance Mathematics (WIM) 2024**, at the **UNSW actuarial seminar**, Australia, and in a seminar at the **Autorités des Marchés Financiers (AMF)**, Québec.

From Uncertainty to Precision: Enhancing Binary Classifier Performance through Calibration

by Agathe Fernandes Machado, Arthur Charpentier, Emmanuel Flachaire, Ewen Gallic, François Hu [arXiv:2402.07790](https://arxiv.org/abs/2402.07790)

The assessment of binary classifier performance traditionally centers on discriminative ability using metrics, such as accuracy. However, these metrics often disregard the model's inherent uncertainty, especially when dealing with sensitive decision-making domains, such as finance or healthcare. Given that model-predicted scores are commonly seen as event probabilities, calibration is crucial for accurate interpretation. In our study, we analyze the sensitivity of various calibration measures to score distortions and introduce a refined metric, the Local Calibration Score. Comparing recalibration methods, we advocate for local regressions, emphasizing their dual role as effective recalibration tools and facilitators of smoother visualizations. We apply these findings in a real-world scenario using Random Forest classifier and regressor to predict credit default while simultaneously measuring calibration during performance optimization.




This work has been presented at the StatQAM seminar, and will be presented this Spring and Summer (**Journées Canadiennes de Sciences Économiques** and **Insurance Data Science** conference).

Geospatial Disparities: A Case Study on Real Estate Prices in Paris

by Agathe Fernandes Machado, François Hu, Philipp Ratz, Ewen Gallic, Arthur Charpentier [arXiv:2401.16197](https://arxiv.org/abs/2401.16197)

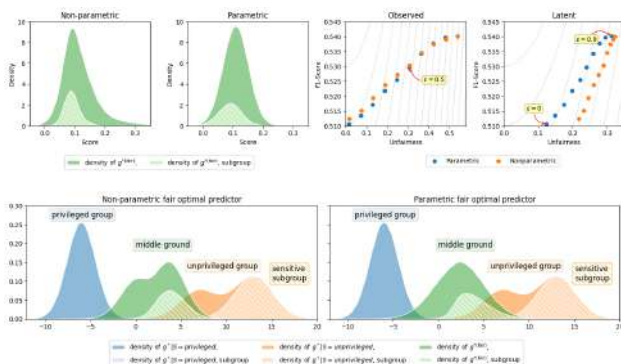
Driven by an increasing prevalence of trackers, ever more IoT sensors, and the declining cost of computing power, geospatial information has come to play a pivotal role in contemporary predictive models. While enhancing prognostic performance, geospatial data also has the potential to perpetuate many historical socio-economic patterns, raising concerns about a resurgence of biases and exclusionary practices, with their disproportionate impacts on society. Addressing this, our paper emphasizes the crucial need to identify and rectify such biases and calibration errors in predictive models, particularly as algorithms become more intricate and less interpretable. The increasing granularity of geospatial information further introduces ethical concerns, as choosing different geographical scales may exacerbate disparities akin to redlining and exclusionary zoning. To address these issues, we propose a toolkit for identifying and mitigating biases arising from geospatial data. Extending classical fairness definitions, we incorporate an ordinal regression case with spatial attributes, deviating from the binary classification focus. This extension allows us to gauge disparities stemming from data aggregation levels and advocates for a less interfering correction approach. Illustrating our methodology using a Parisian real estate dataset, we showcase practical applications and scrutinize the implications of choosing geographical aggregation levels for fairness and calibration measures.



Parametric Fairness with Statistical Guarantees

by François Hu, Philipp Ratz, Arthur Charpentier [arXiv:2310.20508](https://arxiv.org/abs/2310.20508)


Algorithmic fairness has gained prominence due to societal and regulatory concerns about biases in Machine Learning models. Common group fairness metrics like Equalized Odds for classification or Demographic Parity for both classification and regression are widely used and a host of computationally advantageous post-processing methods have been developed around them. However, these metrics often limit users from incorporating domain knowledge. Despite meeting traditional fairness criteria, they can obscure issues related to intersectional fairness and even replicate unwanted intra-group biases in the resulting fair solution. To avoid this narrow perspective, we extend the concept of Demographic Parity to incorporate distributional properties in the predictions, allowing expert knowledge to be used in the fair solution. We illustrate the use of this new metric through a practical example of wages, and develop a parametric method that efficiently addresses practical challenges like limited training data and constraints on total spending, offering a robust solution for real-life applications.




Addressing Fairness and Explainability in Image Classification Using Optimal Transport

by François Hu, Philipp Ratz, Arthur Charpentier [arXiv:2308.11090](https://arxiv.org/abs/2308.11090)

Ensuring trust and accountability in Artificial Intelligence systems demands explainability of its outcomes. Despite significant progress in Explainable AI, human biases still taint a substantial portion of its training data, raising concerns about unfairness or discriminatory tendencies. Current approaches in the field of Algorithmic Fairness focus on mitigating such biases in the outcomes of a model, but few attempts have been made to try to explain why a model is biased. To bridge this gap between the two fields, we propose a comprehensive approach that uses optimal transport theory to uncover the causes of discrimination in Machine Learning applications, with a particular emphasis on image classification. We leverage Wasserstein barycenters to achieve fair predictions and introduce an extension to pinpoint bias-associated regions. This allows us to derive a cohesive system which uses the enforced fairness to measure each features influence on the bias. Taking advantage of this interplay of enforcing and explaining fairness, our method hold significant implications for the development of trustworthy and unbiased AI systems, fostering transparency, accountability, and fairness in critical decision-making scenarios across diverse domains.




Artificial Intelligence and Personalization of Insurance: Failure or Delayed Ignition?

by Xavier Vamparys, Arthur Charpentier, 


In insurance, there is still a significant gap between the anticipated disruption, due to big data and machine learning algorithms, and the actual implementation of behavior-based personalization, as described by Meyers (2018). Here, we identify eight key factors that serve as fundamental obstacles to the radical transformation of insurance guarantees, aiming to closely align them with the risk profile of each policyholder. These obstacles include the collective nature of insurance, the entrenched beliefs of some insurance companies, challenges related to data collection and use for personalized pricing, limited interest from insurers in adopting new models, as well as policyholders' reluctance towards embracing connected devices. Additionally, the hurdles of explainability, insurer inertia, and ethical or societal considerations further complicate the path toward achieving highly individualized insurance pricing.

A version in French appeared in the *Cahiers de la Chaire PARI* collection (# 32).


Following that publication, Arthur Charpentier had a live twitch interview for  [la lettre de l'assurance](#) (in French).



Melting contestation: insurance fairness and machine learning

by Laurence Barry, Arthur Charpentier,  doi: [10.1007/s10676-023-09720](https://doi.org/10.1007/s10676-023-09720), in *Ethics and Information Technology* journal.

With their intensive use of data to classify and price risk, insurers have often been confronted with data-related issues of fairness and discrimination. This paper provides a comparative review of discrimination issues raised by traditional statistics versus machine learning in the context of insurance. We first examine historical contestations of insurance classification, showing that it was organized along three types of bias: pure stereotypes, non-causal correlations, or causal effects that a society chooses to protect against, are thus the main sources of dispute. The lens of this typology then allows us to look anew at the potential biases in insurance pricing implied by big data and machine learning, showing that despite utopic claims, social stereotypes continue to plague data, thus threaten to unconsciously reproduce these discriminations in insurance. To counter these effects, algorithmic fairness attempts to define mathematical indicators of non-bias. We argue that this may prove insufficient, since as it assumes the existence of specific protected groups, which could only be made visible through public debate and contestation. These are less likely if the right to explanation is realized through personalized algorithms, which could reinforce the individualized perception of the social that blocks rather than encourages collective mobilization.




A Sequentially Fair Mechanism for Multiple Sensitive Attributes

by François Hu, Philipp Ratz, Arthur Charpentier, [arXiv:2309.06627](https://arxiv.org/abs/2309.06627)

In the standard use case of Algorithmic Fairness, the goal is to eliminate the relationship between a sensitive variable and a corresponding score. Throughout recent years, the scientific community has developed a host of definitions and tools to solve this task, which work well in many practical applications. However, the applicability and effectivity of these tools and definitions becomes less straightforward in the case of multiple sensitive attributes. To tackle this issue, we propose a sequential framework, which allows to progressively achieve fairness across a set of sensitive features. We accomplish this by leveraging multi-marginal Wasserstein barycenters, which extends the standard notion of Strong Demographic Parity to the case with multiple sensitive characteristics. This method also provides a closed-form solution for the optimal, sequentially fair predictor, permitting a clear interpretation of inter-sensitive feature correlations. Our approach seamlessly extends to approximate fairness, enveloping a framework accommodating the trade-off between risk and unfairness. This extension permits a targeted prioritization of fairness improvements for a specific attribute within a set of sensitive attributes, allowing for a case specific adaptation. A data-driven estimation procedure for the derived solution is developed, and comprehensive numerical experiments are conducted on both synthetic and real datasets. Our empirical findings decisively underscore the practical efficacy of our post-processing approach in fostering fair decision-making.

See also [equipy](#) package.



Est-il nécessaire (et utile) d’être en guerre contre tout ? (Dissemination)

by Arthur Charpentier, published in *Risques*.

ÉTUDES ET DÉBATS | ARTHUR CHARPENTIER

Est-il nécessaire (et utile) d’être en guerre contre tout ?

Début 2010, Nicolas Sarkozy, suivi par Jérôme Cahuzac, entraînait « en guerre contre les fraudeurs ». En mars 2020, la France (par la voix de son président) entraînait en « guerre sanitaire » contre un coronavirus. Et plus récemment, le sociologue Dominique Ména reprenait l’expression de l’économiste Christian Gollier qui commençait son livre par la phrase : « dans ce livre, j’exprime mes espoirs et mes doutes quant à la possibilité de gagner la guerre mondiale contre les dérèglements du climat ». Évoquer la guerre permet probablement de marquer les esprits, mais à force d’être en guerre contre tout on peut se demander si l’association a encore du sens, surtout dans un contexte de gestion des risques.



ARTHUR CHARPENTIER
Professeur, UQAM
(Université du Québec à Montréal)

Le vocabulaire guerrier en politique (et en journalisme)

Les déclarations de guerre symboliques ont commencé aux États-Unis dans les années 1960, avec la « guerre contre la pauvreté » puis contre « le crime » par l’administration de Lyndon Johnson en 1964 et 1965, puis contre « la drogue » et contre « le cancer » en 1971, par l’administration de Richard Nixon. À quoi servent ces déclarations de guerre ? « Nous devons employer le langage utilisé en temps de guerre pour donner le ton juste à ce que nous vivons. L’urgence doit de créer de la confusion au sein de la population », a écrit Antoine-Armand Dailiac, le général qui a commandé les forces des Nations unies au Rwanda lors du génocide. « Ces termes doivent être utilisés pour permettre aux gens de comprendre l’urgence de la situation et l’engagement essentiel dans une campagne offensive. Mais on ne voit, ce langage guerrier est un choix politique comme le rappelle Dailiac (2020) : « il s’agit de faire peur, de pousser la population à observer un comportement d’urgence. Attention, je ne dis pas que la situation n’est pas grave. Je dis que les mots employés ne sont pas dictés par elle mais par la stratégie des pouvoirs publics pour tenter de résoudre la crise. »

Quelques guerres en 2016 et 2017

2017
GUERRE CONTRE LE PLASTIQUE
“War on ocean plastic”
UNEP

GUERRE CONTRE L’ÉPICERIE
“Grocery war”
Radford

GUERRE CONTRE LE GASPILLAGE
“War on waste”
Liu

2016
GUERRE CONTRE LES JEUX VIDÉO
“War on video games”
Tassi

Partage des données, à qui profite le crime ? (Dissemination)

by Arthur Charpentier, Laurence Barry, published in *Risques* (forthcoming).

Manuel d’Assurance

A revised version of the *Manuel d’Assurance* is now available, as e-book, since January 2024. Arthur Charpentier gave some interviews ([▶](#)).



Manuel d'assurance

Gilles Bénéplanc Arthur Charpentier Patrick Thourot

Collection : **Savoirs**


Discipline : **Droit et Science politique**

Catégorie : **Livre numérique**

Date de parution : **02/11/2022**

📅 Organisation of Conferences

Organization of a workshop **Networks, Games and Risk**, december 2023, at UQAM, with Renaud Bourles (Centrale Marseille, Aix-Marseille School of Economics), Vincent Boucher (Université Laval), Federico Bobbio (Université de Montréal), Leonie Baumann (McGill University), Fallou Niakh (CREST-ENSAE, Institut Polytechnique de Paris) and Philipp Ratz (UQAM), 🌐.



Fallou Niakh (PhD student CREST-ENSAE) was invited for a month at UQAM.

📅 Presentations in Conferences

AAAI'24, Vancouver

François presented “A Sequentially Fair Mechanism for Multiple Sensitive Attributes” at the AAAI'24 conference, 38th Annual AAAI Conference on Artificial Intelligence, 🌐.

A Sequentially Fair Mechanism for Multiple Sensitive Attributes
François H.C., Philipp Ratz¹ and Arthur Charpentier²
¹Université de Montréal (francois.h.c. @umontreal.ca), ²Université de Québec à Montréal

MOTIVATION: INTERSECTIONAL FAIRNESS
Insurance regulations affect various features and always across some latent feature. Consider the following example: (1) a person with a certain age and gender is charged a higher premium than a person with a different age and gender. (2) a person with a certain age and gender is charged a higher premium than a person with a different age and gender. (3) a person with a certain age and gender is charged a higher premium than a person with a different age and gender.


SEQUENTIAL FAIRNESS MECHANISM
The core of the problem is to find a sequentially fair mechanism for multiple sensitive attributes. This is achieved by a sequentially fair mechanism for multiple sensitive attributes. This is achieved by a sequentially fair mechanism for multiple sensitive attributes.

NUMERICAL RESULTS
Table 1: In the table, we report the results of the numerical experiments. The table shows the results of the numerical experiments. The table shows the results of the numerical experiments.

BACKGROUND: WASSERSTEIN BARYCENTER
The Wasserstein barycenter is a generalization of the Wasserstein distance. It is defined as the Wasserstein barycenter of a set of probability distributions. It is defined as the Wasserstein barycenter of a set of probability distributions.

EXTENSION TO APPROXIMATE FAIRNESS
In a nutshell (Figure 4), introducing a sequential process is phased to be on taking fairly. The mechanism is extended to approximate fairness, which involves approximating the fairness constraints with a given level of fairness approximation using fairness. The study has different components to achieve this.

REFERENCES
[1] ... [2] ... [3] ... [4] ... [5] ... [6] ... [7] ... [8] ... [9] ... [10] ...



Videos are online, 📺 (and bloopers 📺)

François presented this work (on fairness with multiple sensitive attributes) also at the 17th Financial Risks International Forum “Big Data & Algorithmic Finance”, in Paris, in March 2024, 🌐.

